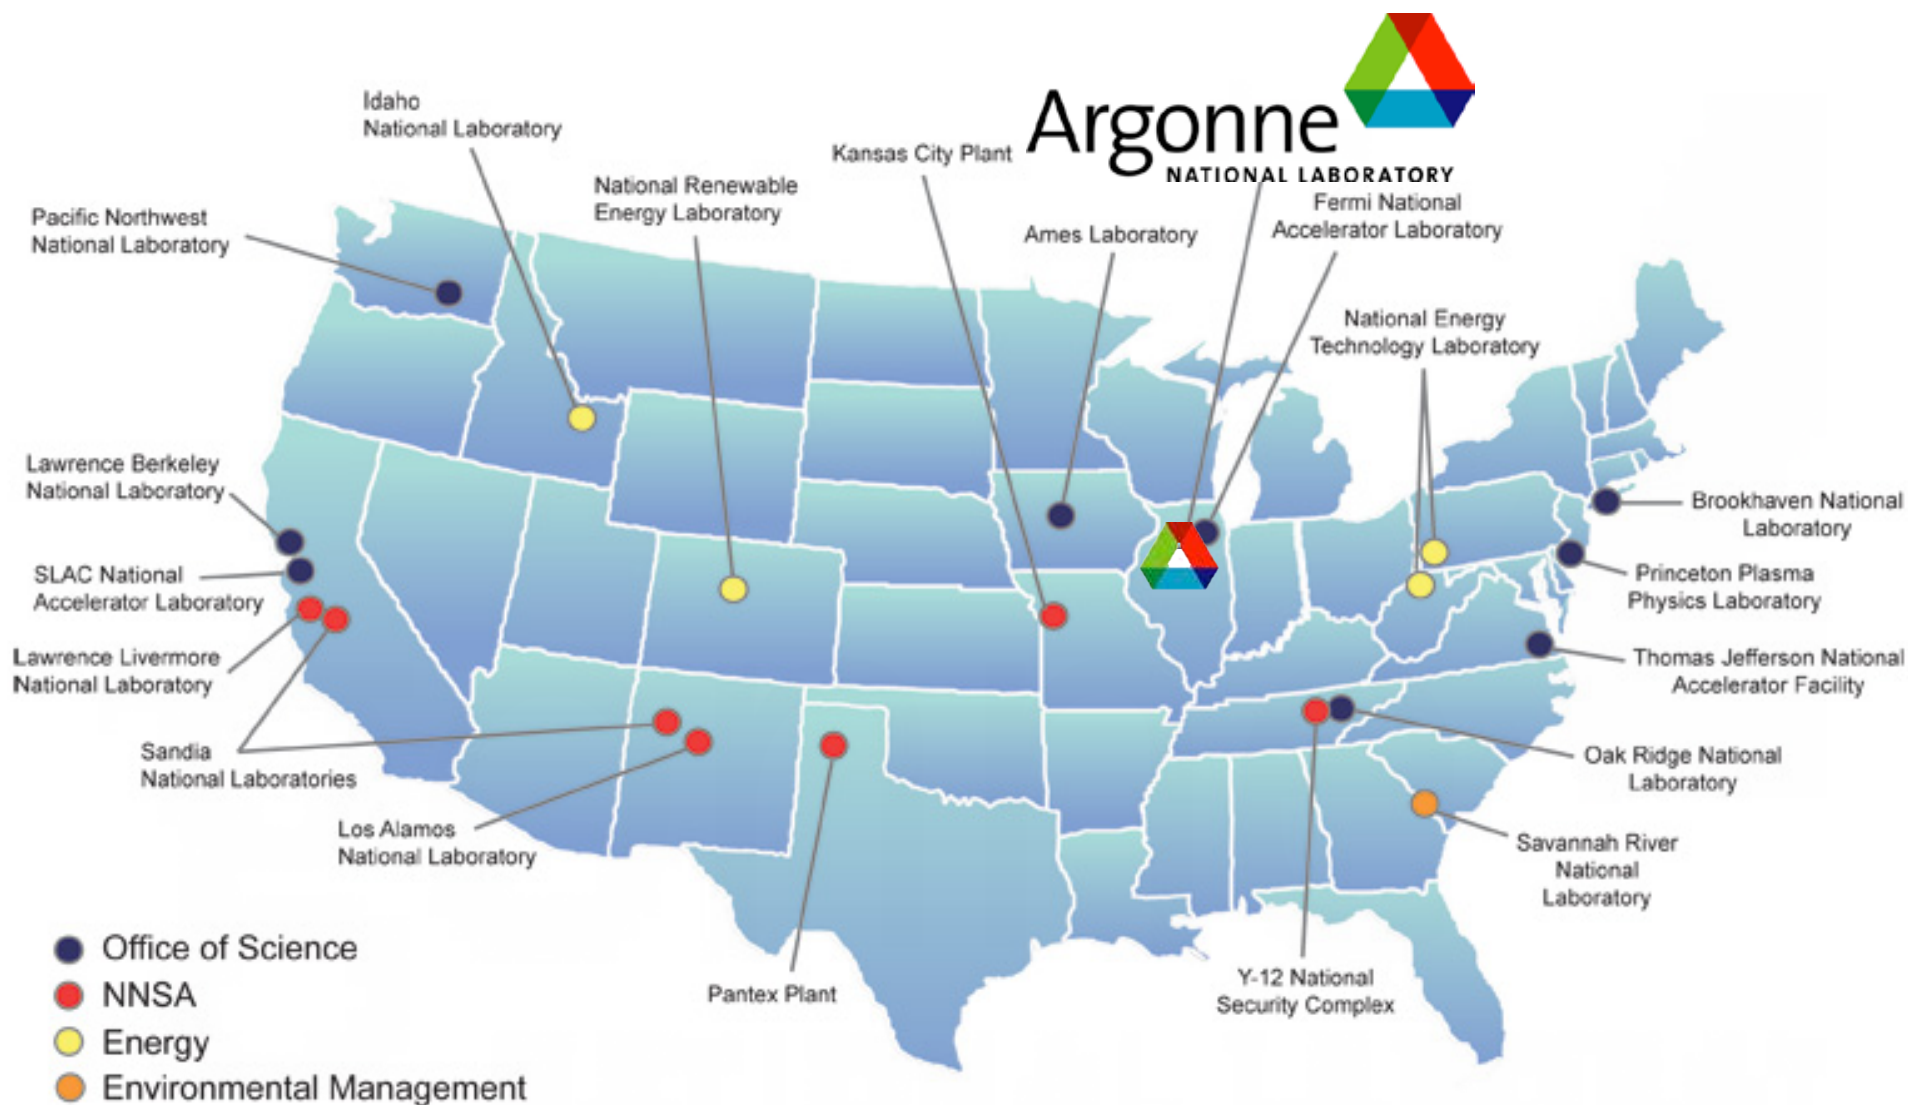


An Introduction to Parallel Supercomputing

Pete Beckman

Argonne National Laboratory

Argonne & the DOE National Laboratory System



Argonne National Laboratory

- \$675M /yr budget
- 3,200 employees
- 1,450 scientists/eng
- 750 Ph.D.s



Direct descendent of Enrico Fermi's Metallurgical Laboratory



- Opened in Feb 1943 (as new site for Chicago's Metallurgical Laboratory)
- Became Argonne National Laboratory in July 1946 (first national laboratory)



MCS Division meeting c. 1983

- “If our R&D is going to be relevant ten years from now, we need to shift our attention to parallel computer architectures”
- “Los Alamos has a Denelcor HEP: let’s experiment with it”





POOMA Project: 1996

John Reynders



Parallel Platform Paradox

“The average time required to implement a moderate-sized application on a parallel computer architecture is equivalent to the half-life of the latest parallel supercomputer.”

“Although a strict definition of “half-life” could be argued, no computational physicist in the fusion community would dispute the fact that most of the time spent implementing parallel simulations was focused on code maintenance, rather than on exploring new physics. Architectures, software environments, and parallel languages came and went, leaving the investment in the new physics code buried with the demise of the latest supercomputer. There had to be a way to preserve that investment.”



Pete's Investment Recommendations

- **Other People's Libraries**
- **Encapsulation**
 - Parallelism & Messaging & I/O
- **Embedded Capabilities**
 - Debugging
 - Performance Monitoring
 - Correctness Detection
 - Resilience
- **The Two Workflow Views**
 - Science: (problem setup, analysis, etc.)
 - Programmer: (mod, testing, document, commit)
- **Automation**
 - A+ Build system, nightly test and build, configuration
 - Embedded versioning and metadata
- **Community: web, tutorial, email, bug tracking, etc**



Encapsulation Examples from Mike Heroux...

dft_fill_wjdc.c

Tramonto
WJDC
Functional

- New functional.
- Bonded systems.
- 552 lines C code.

WJDC-DFT (Wierthm, Jain, Dominik, and Chapman) theory for bonded systems. (S. Jain, A. Dominik, and W.G. Chapman. *Modified interfacial statistical associating fluid theory: A perturbation density functional theory for inhomogeneous complex fluids. J. Chem. Phys.*, 127:244994, 2007.) Models stoichiometry constraints inherent to bonded systems.

How much MPI-specific code?

dft_fill_wjdc.c
MPI-specific
code

source_dp_g.f

MFIX
Source term for
pressure
correction

- MPI-callable, OpenMP-enabled.
- 340 Fortran lines.
- No MPI-specific code.
- Ubiquitous OpenMP markup (red regions).



Threads/Tasks: Managing Exploding Parallelism

■ Dynamic parallelism and decomposition

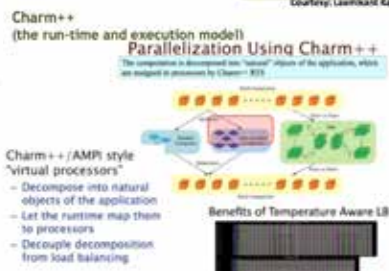
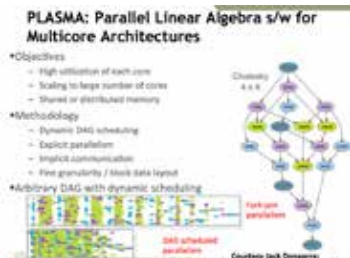
- Programmer cannot hand-pick granularity / resource mapping
 - (equal work != equal time)



≠




Variability is the new norm:
Power
Resilience
Intranode Contention



Fault-Tolerance is Already Here

- We have already seen the future
 - 8 years ago in fact.
- Persistent ECC memory faults are the norm, not the exception
 - Machines need to stay up to satisfy their contracts.
 - Over the course of a day or two these parts can be replaced, but not over the life of your batch job.



From Brian Van

Straalen

Future Trends: Invest Wisely

Trending Up

Trending Down

| | |
|-------------------------------------|--|
| Asynchrony, Latency Hiding | Block synchronous |
| Over Decomp & Load Balancing | Static partitioning per core |
| Massive Parallelism | Countable parallelism |
| Reduced RAM per Flop | Whole-socket shared memory |
| Software-managed memory | Simple NUMA |
| Expensive Data Movement | Expensive flops |
| Fault / Resilience Strategies | Pure checkpoint/restart |
| Low BW to Storage, in-situ analysis | Save all, let the viz guys sort it out |



Of Refractions.

If y^e ray be refracted at the center of y^e circle acc^d towards d & $ab \perp be \perp gc$ $|| d$.
Then suppose $ab : ed :: d : e$. Q.E.D.

Then suppose $ac \parallel \delta \epsilon$. See Cartes Dioptricks
If there be an hyperbola whose foci a and c
are to its transverse axis bf as δ to ϵ . Then y^e ray ac $\parallel \delta \epsilon$
is refracted to y^e exterior focus d . See C. Dioptricks

3 Having y^2 proportion of d to z , or, $bd:bf$.
The Hyperbola may be thus described.

Upon y^e centers a. b.
let y^e instrument adbe
be moved in such instrument
observe y^e ad \perp de \perp c f
of y^e the beam cut is
not in y^e same plane as
adbe but intersects it at
 y^e angle $\angle ev$ soe y^e if
 $tv \perp ev$, then D:ex: \angle f:tv.
Or D:e :: Rud: sine of \angle ev.
Also make $de = \frac{y}{2}$ i.e
half y^e transverse
the figure diam

Then place ^{the finished diamond} y^e plate ^{on the} ~~in~~ in the same plane with ab
 & moving y^e instrument adbest to & fro its edge
 cet shall cut or wear it ~~to~~ into y^e shape of y^e
 desired Parabola. Or the plate ~~can~~ may be filed
 away untill y^e edg cet exactly touch it every where.

2 By the same proceeding D_{45} = Caries concave
Hyperbolicall wheels may be described by being
turned with a chissell D_{46} whose edge is a straight
line inclined to the ~~axis~~ of the mandrill by a
well angle is found by making D_{47} = D_{46} by g = h

3. By the same reason a whole may be turned Hyperbolically concave, if Hyperbola being convex. Or a Plate may be turned Hyperbolically concave.

